

Fault Management Metrics

Stephen B. Johnson¹

Dependable System Technologies, LLC, Broomfield, CO, 80020

Sudipto Ghoshal²

Qualtech Systems, Inc., Rocky Hill, Connecticut, 06067

Deepak Haste³

Qualtech Systems, Inc., Rocky Hill, Connecticut, 06067

and

Craig Moore⁴

NASA Marshall Space Flight Center, Huntsville, AL, 35812

This paper describes the theory and considerations in the application of metrics to measure the effectiveness of fault management. Fault management refers here to the operational aspect of system health management, and as such is considered as a meta-control loop that operates to preserve or maximize the system's ability to achieve its goals in the face of current or prospective failure. As a suite of control loops, the metrics to estimate and measure the effectiveness of fault management are similar to those of classical control loops in being divided into two major classes: state estimation, and state control. State estimation metrics can be classified into lower-level subdivisions for detection coverage, detection effectiveness, fault isolation and fault identification (diagnostics), and failure prognosis. State control metrics can be classified into response determination effectiveness and response effectiveness. These metrics are applied to each and every fault management control loop in the system, for each failure to which they apply, and probabilistically summed to determine the effectiveness of these fault management control loops to preserve the relevant system goals that they are intended to protect.

I. Introduction

FAULT Fault Management (FM) is the operational aspect of System Health Management (SHM), which in turn is defined as “the capabilities of a system that preserve the system's ability to function as intended”.¹ Alternatively, SHM can be defined as the capabilities of a system that preserve the system's ability to achieve its goals. Fault Management consists of the non-passive aspects of these capabilities that detect and respond to projected or existing failure within the system. These detection and response capabilities function as “meta-control loops” that determine if the system's nominal control functions are operating properly, and if not, to place the system into a state that is once again controllable. This state could be controlled either to existing or degraded system goals.²

Historically, one of the major difficulties in designing FM has been measuring or determining the value of differing FM designs. To date, the vast majority of FM design has been performed on an ad hoc basis, with some use of qualitative requirements such as single failure tolerance, assessed against a suite of failure modes provided in failure modes and effects analyses (FMEA). Quantitative estimates of the value of FM have occasionally been performed on an ad hoc basis to assess specific design options.

¹ President, and AIAA Member.

² Vice President, Engineering, not AIAA Member.

³ Director, Engineering, not AIAA Member.

⁴ Aerospace Engineer, Integrated System Health Management and Automation Branch, EV43, not AIAA Member.

As in any other engineering design, FM design should be based on knowledge of the value it provides. Since for a given system there could be a variety of potential FM designs, how does a designer choose which is best? When has enough FM been put into the system? How much is “too much”?

Given that FM functions as a suite of control loops, heretofore called “Fault Management Control Loops” (FMCL), the question of value for FM boils down to the question of how much value each possible FMCL provides, and collectively, how much they cumulatively provide for the system. Having quantitative measures of this value is preferable so as to enable trade studies of different design options. Quantitative measurement or estimate of value implies the need for metrics.

This paper provides an overview of FM metrics, based on the above-noted theory that FM operates as meta-control loops. Given this theory, this paper proposes and demonstrates that the value of FM can be measured quantitatively in a manner similar to, but extended beyond the measures used to assess the value of classical feedback control systems.

II. Extending Control Theory

The idea that FM is an extension of control theory implies that the key FM metrics are divided into metrics for state estimation (detection, prognostics, and diagnostics—isolation and identification), and for control (response decision, and response action).³ These two general categories of state estimation and control relate directly to two distinct categories of metrics: metrics for state estimation, and metrics for state control. These ultimately combine to form a single metric related to the probability of achieving the goal of the FMCL. That is, the probability of achieving the goal of the FMCL is the measure of how well the FMCL provides “goal insurance.”

In general, the estimate per FMCL is a set of contingent probabilities, starting with the coverage estimates of detection capabilities. For those failures (and/or degradations and/or anomalies) that are detected, then the next step is to calculate the effectiveness of the detection. If this is successful, then the effectiveness of the diagnostic capabilities are assessed. For that fraction of successful diagnoses, the system’s ability to make the correct response decision is estimated. For that fraction for which the correct decision is made, then the response must be successfully performed, and its capability to do so estimated. This resulting contingent probability is the final value of the effectiveness of the total FMCL.

State estimation metrics always relate to the probability of determining the “true” state of the system. These are often described in “truth tables” (or “confusion matrices”) typically represented as 2x2 matrices. The cells in the matrix relate the relationship between the true state and the estimate of the true state. For failure detection, these four relationships are listed and defined as follows, with respect to the determination as to whether a failure exists.

- 1) True Positive (TP): A correct estimate that a failure exists, when it really does exist.
- 2) True Negative (TN): A correct estimate that a failure does not exist, when it really does not exist.
- 3) False Positive (FP): An incorrect estimate that a failure exists, when it really does not exist.
- 4) False Negative (FN): An incorrect estimate that a failure does not exist, when it really does exist.

Each of these four is estimated in probabilistic terms, and those estimates in turn have uncertainties associated with those probability estimates. State estimation functions are divided into three major types: detection, diagnostics, and prognostics. The FP/FN/TP/TN definitions differ slightly for diagnostics and prognostics, as compared to detection, which is defined above.

State control metrics for decision and response are estimated based on their effectiveness in selecting and executing the proper response. As with control in general, these responses must execute in a timely manner before the associated system goal of the FMCL is compromised. The timeliness depends not only on the response, but on the physics of the failure effects that are being responded to, and the latencies of the entire FMCL, including the detection and diagnostic functions as well as the decision and response latencies.

III. Failure Scenarios

In general, any given FMCL detects and responds to a specific set of failure behaviors, which typically result from several or many failure modes. That is, many failure modes ultimately produce the same “intermediate failure effect” that is detected by the failure detection mechanism of an FMCL. While it is typical for FMEAs to define failure modes and then determine their effects on the system, from the perspective of FM, it is important to group failure modes by those common effects that are detected by a particular FMCL or group of FMCLs. For example, on a satellite one might have a failure detection that monitors the rotation rates of a vehicle. During attitude (rotation) maneuvers, the system expects certain rates to occur and will flag a failure if those rates go too high or too low. Let’s assume that the rates are too low. This could be because a thruster valve is stuck closed, because there is debris in a propellant line, or an isolation valve failed to a closed position. In turn, these problems could be caused by a

mechanical problem in a valve, a power failure to a valve, spurious commands to an isolation valve, or a host of other problems. However all of these problems are picked up by a single detection mechanism. To perform the calculation of detection effectiveness for the attitude rate detection “low”, one must group together all possible causes (failure modes) in all possible vehicle configurations and system modes and phases, as in some modes the failure may not be detectable, and in others it is easy to detect. To determine the value of the attitude rate detection and response which would be to switch to an alternate set of thrusters (the FMCL), one must also add up the probabilities of all of the failure modes that can cause the detection to occur, as the benefit of the detection is based on its ability to mitigate the set of failures that it can successfully detect and mitigate (respond to).

As noted in this example, calculations of metrics for individual FMCLs requires consideration of major differences in failure behaviors and probabilities based on mission phases, modes and system configurations. Differing failure modes produce different system failure effects that are detected by different FMCLs. Furthermore, even for the same failure mode or group of failure modes that produce a given effect in one vehicle configuration, the end-effects can vary based on the system configuration and the external environment if it changes over time.

To address these variations, the FM analyst must produce a global set of “failure scenarios”, each differentiated by the capability of the FMCLs to detect and respond, and to relevant system and environmental changes that alter the effectiveness of the FMCL. In general, because many failure modes produce the same failure effects at a higher level, and because these higher-level effects are what are detected and responded to by FMCLs, there are fewer failure scenarios than there are failure modes. For example, on the National Aeronautics and Space Administration’s (NASA) Space Launch System (SLS) there are already over 16,000 identified failure modes, but overall these yield about 1,300 failure scenarios that can create loss of mission and hence cause threats to the safety of the crew. Each must be analyzed, though there are significant similarities among many of them, so that this analysis is further reduced to something over 100 failure cases that vary in results across mission phases. If the analysis is properly separated, the phase-dependent effects based on major changes to the external environment and system configuration can be analyzed separately, and then combined into the 100 specific cases.

The FM analyst must specify all specific failure scenarios, and globally, define the entire suite of failure scenarios. Ideally, these are defined as a set of failure modes in a particular system configuration and mission phase and mode, which if not mitigated will compromise a relevant system goal. A comprehensive set of scenarios would encompass all failure modes across all relevant system configurations and mission phases/modes. For each defined scenario thus defined would be a set of FM metrics for the individual pieces of the FMCL, which are then combined in each FMCL as the product of contingent probabilities against relevant goals, and failures at each step contributing to estimates that the goal is not met in that scenario. Finally, the estimates from all scenarios are summed, probabilistically weighted via the probability either of the failure modes, or of the scenario as a whole independent of the failure modes.

IV. FM Metrics Theory and Description

This section describes each of the major FM metrics in turn, from detection coverage, to detection effectiveness, to diagnostic and/or prognostic effectiveness, to response determination, to response effectiveness. These individual metrics are then combined to yield the effectiveness metrics for FMCLs, and from there, to the effectiveness of the FM for the system as a whole.

A. Failure Detection Coverage

The first major metric related to detection is coverage. A detection mechanism can be assessed for effectiveness only against those failures, degradations, or anomalies that it theoretically capable of detecting. If it cannot detect these behaviors, even in theory, then the detection has no “coverage” of those behaviors. Conversely, if behaviors are potentially detectable by the detection mechanism, then the mechanism “covers” them.

Both failures and degradations relate to factors of “unacceptability”, whereas anomalies relate to “unexpectedness”. However, all of them ultimately relate to achievement of system goals. A more detailed discussion of the nature of goals the their relationship to failure and to failure detections can be found in Reference 4.⁴

Since coverage must be assessed against a goal, one must specify the goal against which the coverage metric is assessed. Once that is specified, then one can determine which failure modes whose failure effects can potentially compromise the achievement of that goal. Then, one identifies the detections available and can determine if they can collectively detect all of these failures. One can specify the fraction of failure modes whose effects can compromise the goal that are detected, divided by the total number of failure modes that can compromise the goal. This fraction is a simple coverage metric. To increase accuracy, one can then also define the probability of each failure mode

occurring, and generate a metric of the total amount of probability that can compromise the goal that is detected, divided by the total amount of probability that the goal can be compromised. This latter metric is the most useful as it is in a probabilistic form that will ultimately be needed. All undetected failure modes and their probabilities directly compromise the goal as no detection or response will occur. All detected failure modes and their probabilities are then passed along to the next step in state estimation calculations.

One key nuance to the coverage estimate is deciding how to classify failure modes whose behaviors sometimes lead to failure and sometimes do not, or when those behaviors, even if they cause failure of some sub-function, do not always cause higher level function failures. This issue helps make clear why it is important to define the goal against which one estimates coverage, as it helps clarify what behaviors “matter” and what behaviors do not. It may be necessary to subdivide the probabilities into covered and non-covered fractions for some failure modes.

Coverage for detection of degradation is estimated in the same way as coverage for failure detection, as degradation detection is performed against the same state variables and the same goals as failure, but with tighter threshold limits. For degradation, thresholds are set closer to “ideal” operation and are still “acceptable” performance, though no longer “ideal”.

Detection of anomalies differs significantly than for failure and degradation, as degradation detection merely detects differences between expected and actual behavior. Expectations in turn can be based on prior analysis, or on prior operation. Without some data that defines expectations, no anomaly detection can occur. In general, one can say that anomaly detection is also based on goals, as anomaly detection picks up differences between “unexpected” performance of functions. Anomalies are defined as unexpected behaviors, which are unexpected values or patterns in state variables, which in turn when constrained are related to goals. While it is conceivable to find anomalies in state variables independently of goals, the only state variables worth checking for anomalies are those related to system goals. Thus in terms of coverage, it is probably best to start by determining how state variables in the system are related to goals, and to check for anomalies of state variables related to goals worthy of the effort. Thus coverage of anomaly detection is ideally based on a comparison of the total number of state variables that are associated with “sufficiently important” goals, and then determining if all of these are addressed by anomaly detection mechanisms. A nuance to this method is the assessment of whether the values of some not-directly-monitored state variables can or should be calculated based on other monitored state variables, and then checking for anomalies based on this calculated state instead of direct observation.

For those anomalies, degradations, and failures that cannot be detected even in theory and hence are “not covered”, all non-covered behaviors are immediately removed from any calculations of the effectiveness of the relevant FMCLs. It is invalid to allocate failure to detect an anomaly, degradation, or failure to an FMCL if it is impossible even in theory for that FMCL to detect it. These “failures to detect” anomalies, degradations, or failures are charged against the system as a whole, but not to any individual FMCL. For example, if a particular failure mode is undetectable, but if it occurs it will cause loss of human life, then the system’s ability to protect against human life must be assessed as lower by the probability of the occurrence of the failure mode. This is more difficult to estimate probabilities in this way for anomalies, because the purpose of anomaly detection is to detect unexpected behaviors, including those not predicted in any prior analysis.

B. Failure Detection Effectiveness

Failure detection is the function of “deciding that a failure exists.” Since one way of defining failures is “the inability to achieve a system goal”, it is clear that when deciding if a failure exists, one must specify the goal that is to be achieved or the function that is to be performed. Then one must decide whether the behaviors being monitored indicate that the goal will be achieved.

For all detected failure behaviors, one calculates the truth table metrics of TP/TN/FP/FN of the various detection mechanisms. Truth table metrics for detection all theoretically use probability distributions of successful and unsuccessful detections. The calculation of truth table probabilities must account for two major factors: the value of the detection threshold, and the probability of failure of the detection mechanisms. An example of the details for such calculations are described in Reference 5.⁵

The first issue relates to the raw physics of the behaviors being monitored, and how the threshold that separates “failure” from “non-failure” is determined. Thresholds can be set closer to nominal behaviors, or further away. If set closer to nominal behaviors, then the probability of successfully detecting failure (TP) improves, but the possibility of incorrectly deciding a failure exists when it really does not (FP) also goes up. Conversely, if the threshold is set far from nominal behaviors, then the probability of TN goes up but FP goes down. However, the probability of not detecting the failure (FN) goes up. Estimation of these probabilities can be based on physics-based analysis prior to system operation, or from actual data from system operation. In practice on Space Launch System, we have found

that it is most crucial to calculate FP and FN, and then $1 - FN = TP$. TN does not generally appear in the calculations, though whether this is something specific to SLS or is a more general phenomenon is not clear.

The second issue is the failure of the detection mechanism itself. The detection mechanism includes all components of the detection, whether the components are hardware, software, or human. The detection mechanism includes the sensing element, such as sensors, humans looking out windows, or software “observing” an internal variable, transport mechanisms for the sensed data, and the decision mechanisms, including software algorithms, thresholds, data qualification routines, or human decisions. Relevant failures thus include failure modes of hardware, software, or humans, with their associated probabilities. These can include failure of sensors, human error is a human is deciding whether a failure exists, an incorrectly loaded threshold value, software failure, and a host of other issues.

Depending on the design of the detection mechanism, failures of the sensing components can produce various end-results in the overall detection, based on the architecture of how these components are connected, such as the levels and kinds of redundancy. The end results also depend on how the information from the detection mechanisms are combined and used, such as voting, averaging, and mechanisms to disqualify bad sensors or hardware.

Within these scenarios, the FN probability appears as a “non-detection” of a system failure, and is hence a contingent probability. That is, you cannot have an FN unless a failure has occurred, so the probability is the probability of the failure times the probability of the failure to detect the failure. The TP probability is the amount of probability that is passed on for further analysis of the FMCL, while an FN is accounted for as a failure of the relevant function.

By contrast, the FP probability shows up as a “cost” of the detection mechanism. Because every new mechanism put into the system has associated unreliability, this unreliability can and does cause problems that did not exist without that mechanism. In the case of failure detections, a FP tells the system that a failure exists in behaviors being monitored by the detection, when in fact the only failure is of the detection mechanism itself. This means that the system will take some response to a non-existent failure. Depending on what the responses are, this can mean loss of a mission, loss of redundancy or margins, loss of science data, and so on. Those would not have occurred except for the fact that the detection exists and that it has some inherent unreliability, which must be factored into any estimate of its value. This shows up in the failure scenarios as a new set of “false positive failure scenarios” that would not have existed if the FMCLs did not exist. So adding Fault Management adds new failure scenarios, which is to be expected because FM adds new mechanisms into the design that themselves can fail.

Failure detection cannot be completely divorced from diagnostics of failures of the detection mechanisms. If Sensor Data Qualification (SDQ) is accounted for as part of the detection mechanism, the SDQ can remove from the failure detection any failed sensors, and thus increase the reliability of the detection and hence decrease the FP values for the detection. The effectiveness of SDQ is based on the ability to diagnose the fact of a sensor failure. The question is how one should account for this in the metrics. This entails a feedback loop from later parts of the FMCL back to the detection mechanism. This can be addressed in design-phase calculations, but is more difficult for a real-time operational system, if such is desired or needed for the application.

On SLS, SDQ is addressed as described in Reference 5.⁵ One of the things desired on SLS was to determine the “value” of SDQ, so what we did was to compare two detection mechanism designs, one with SDQ and one without, so as to determine how much benefit SDQ provided and hence to determine whether certain SDQ algorithms should be added or included in the abort trigger (failure detection) design. This was done successfully on SLS, and the team concluded that some of the SDQ algorithms have high value, while others have less.

C. Degradation Detection Effectiveness

Failure detection is the function of “deciding that a failure exists.” Since one way of defining failures is “the inability to achieve a system goal”, it is clear that when deciding if a failure exists, one must specify the goal that is to be achieved or the function that is to be performed. Then one must decide whether the behaviors being monitored indicate that the goal will be achieved.

Degradation detection metrics, if needed, can be performed in the same way as failure detections, as these operate on the same state variables and in essence provide “early detection” of prospective failures.

D. Anomaly Detection Effectiveness

Anomaly degradation metrics differ from failure detection metrics in a manner similar to how anomalies differ from failures. While some anomalies occur due to failures, other anomalies are not failures at all, but are nonetheless unexpected and hence legitimate anomalies. One can argue that if known failure modes and effects create off-nominal behaviors, that these are by definition not anomalous because they are known potential behaviors. From an analysis viewpoint, they exist in failure models of the system and are not really anomalous. If one occurs in

operations, it may be temporarily anomalous since failures sometimes occur randomly or “unexpectedly”. When the following investigation then determines that the behaviors are based on a known failure mode, the new behavior is no longer considered anomalous. Thus one can argue that to the extent anomaly detection is detecting known failure modes that can be pre-analyzed, it operates on the same set of failure modes as failure detection, but with a different mechanism to detect them and hence a different effectiveness.

However, anomaly detection generally would not exist if all it did was detect known failure modes, because direct “failure detection” mechanisms are generally more effective in detecting known, pre-assessable failures. The primary purpose of anomaly detection is to detect non-predicted, unknown failures and failure modes (causes). By definition, this means that for its primary function, anomaly detection cannot be assessed against known failure modes. This poses difficulties for analysis.

On what basis can anomaly detection be assessed for effectiveness? One possibility is to assess each anomaly detection based on what is “mathematically possible” as opposed to what is “physically possible”. An anomaly detection compares a set of expected behaviors (states over time) of a set of state variables to current operational values of that same set of state variables. To test this ahead of time, feed the anomaly detection the range of mathematically possible states (and over time, behaviors) and compare how well the anomaly detection picks up these ranges of data. Also, since particular state variables and their states are associated with specific system goals, it should be possible to associate anomalies related to those state variables to those goals to the exclusion of other goals. Thus it is possible to classify anomalies by association with relevant system goals, and therefore to analyze them in smaller, more easily assessable groupings than the entire set of state variables for the system. Not knowing the probabilities associated with the possible states, any anomaly detection effectiveness estimate must be based solely on the numbers of state variables and their mathematical ranges.

E. Diagnostic Effectiveness

The next functions in an FMCL are associated with diagnostics: isolating the location of the cause, and identifying the cause of the detected failure behaviors. Because diagnostics is a state estimation function, it too is estimated with truth/confusion table metrics of FP/FN and TP/TN. Fault isolation, which means determining the location of the cause of the detected failure behaviors, is generally assessed with respect to “ambiguity groups”, which define uncertainties in the possible location of the cause of a failure.

We must translate ambiguity group metrics into FP/FN, and TP/TN probabilistic metrics. For fault isolation, in an ambiguity group of three, if all three locations are equally likely, then the probability of correct isolation on the first attempt is 33%. There may be physical or logical criteria that could change the likelihood of one or other of the components being the correct one. If one has to choose one and cannot choose others, then the probability of the correct choice, the True Positive, is only 33%. The probability of an incorrect choice can be considered a False Positive, which would then have a probability of 67%. On the other hand, if it is possible and acceptable to check all three locations, then the TP probability will be 100%, and the FP probability is 0%, if there are no failures of the fault isolation mechanisms themselves. It is also possible that the fault isolation routine may not identify any possible locations, when in fact there must be one or more. This would be a False Negative. In a crime investigation analogy, the police know there “is” a criminal, but have no idea where he might be, and thus cannot take any actions. In sum, to determine the proper truth table metrics, one must understand the operational use of the diagnostic information to translate from ambiguity group metrics to a FP/FN metrics for the FMCL.

Similar logic applies to fault identification for identifying the correct failure modes.

A fault isolation FP generally results in making the wrong decision about what response to take. Thus the correct response is not taken and the failure effects are not mitigated and continue to propagate. Just like a FP detection, the incorrect response and lack of correct response results in loss of redundancy, loss of mission, loss of science, or some other undesirable result. For fault identification, selecting the incorrect cause of the failure has costs in that the wrong responses will be taken. Since knowledge of the failure mode itself often results in changes to future plans for the system, how to repair a component, how to redesign a component so that failure mode becomes less likely, etc., this too has costs for the system.

Fault isolation FN results generally mean that NO response action is taken, because there is no basis to decide what to do since the location of the cause cannot be identified. The failure effects propagate and nothing is done about it, which leads to the same bad results as if no FMCL existed. One might argue that if fault isolation cannot determine what to do, that it is possible for some systems it may be possible to take that system out of service to prevent an even worse effect. Thus with a nuclear power plant, if failure effects exist but the cause cannot be identified, one might shut down the power plant instead of assuming the risk of explosion. This is a valid point, and can be accounted for in the metric calculations as a case that leads to loss of service, though not a threat to life or

property. One must always determine the impact on system responses and goals, if any FM function fails. This impact can lead to variations in the way the overall FM calculations are performed.

F. Prognostic Effectiveness

Prognostics is another state estimation function, which uses “degradation detection” information as a basis to predict future behavior in a set of state variables, and in particular to predict when failure of a function will occur. Ideally prognosticating into the future requires a physics-based model, but when that is not available, sometimes a simple linear assumption is used to predict failure, in which case a data-driven approach can be used. A good overview of prognostics and associated metrics is described in Reference 6.⁶

Prognostic experts generally agree that a primary output of prognostic functions is the prediction of Remaining Useful Life (RUL), with appropriate uncertainty distributions to account for various factors that can influence the prediction. A component or system remains “useful” as long as it is performing the intended or needed function in an acceptable way. When the component is no longer performing acceptably, then it is considered “failed”, as per typical definitions of failure. This point in time is called the End of Life (EOL). RUL is an estimate of how much time is available for an action to be taken to mitigate the failure or reduce its effects. As such it is essentially an estimate of Time to Criticality (TTC).

Other supporting or related prognostic metrics are related to the quality of this estimate. Some classify these related metrics into estimates of accuracy, precision, and convergence. Accuracy refers to the “degree of closeness of a predictive estimate to its actual value.” Precision relates to the variability of predictions. Convergence refers to the degree of improvement in the estimate over time, as EOL approaches.

From the standpoint of Fault Management theory, prognostics is simply another technique of performing state estimates, with the difference that it is an estimate of a future state, as opposed to a past or current state. That being the case, like any other state estimate, it needs to be converted into truth table metrics of False Positive and False Negative. However, prognostics also adds an additional element, which is that its estimate of RUL inherently relates to the required time in which the overall FMCL must operate, which is the TTC.

To date, the specific methods by which RUL, accuracy, precision, and convergence metrics can be converted to FP/FN and TTC have not been developed in detail. However, some directions for how this can be done are apparent. First, the RUL estimate itself is a TTC estimate, and can be used as such and compared with the overall time of FMCL execution to impact the FMCL effectiveness metric as described below. An RUL estimate that is too conservative predicts EOL at an earlier point in time than the “actual” EOL time. In terms of the overall FMCL and system impact, this leads to a system response taken when not (yet) needed, which acts much like an FP. Every FP leads to an unnecessary response. In this case, the response might eventually be necessary, but if taken early, it will lead to added system costs if a component is replaced too early, or if the system is taken out of service when it did not need to be. Conversely, an RUL that is too optimistic is much like an FN, in that the system could fail before the appropriate failure mitigation is taken. In this case, the full consequence of the failure is suffered, which occurs as if the prognostic did not exist. If there are no other FMCLs to address the consequences of the failure, then system failure ensues. If other FMCLs exist to address failure of the component, then the FN of the prognostic will lead to execution of the real-time FMCL, with its corresponding effectiveness. Presumably the prognostic FMCL would not exist unless the resulting failure leads to some negative consequence that cannot be adequately addressed by a real-time FMCL (classical FDIR).

G. Response-Action Determination Effectiveness

Once state estimation functions are complete, a decision must be made about what to do. To our knowledge, there is no “theory” about what metric to apply to this, but clearly the metric must ultimately relate to whether the correct response is selected, given a valid set of state estimation information upon which to make the decision. Obviously, if the state estimates are wrong, then the response is also likely to be incorrect, but this would be attributed to the state estimation functions, not the decision function. So the decision metric is one that specifies the probability of correct or incorrect decision about what action to take. This is less complex if a predetermined set of actions are available, but far more complex if those are not known ahead of time. If the wrong response is selected, then it will not just be ineffective, but it will likely cause loss of other system resources related to that response (such as removing a redundant component from use, even if that component is functioning properly), while the failure continues to propagate to its ultimate bad effect.

H. Response Effectiveness

The last major metric for FM is the effectiveness of the failure response. This metric applies only if the prior state estimations identified the correct system state, and the decision function has selected the correct response.

Response effectiveness, like all other FM metrics, is a probabilistic quantity regarding the likelihood that the response will mitigate the failure effects that it is intended to address in a specific failure scenario. Since the entire FMCL is intended to preserve some goal, the effectiveness of the response is calculated with respect to that goal. For example, if the goal is to maximize the acquisition of science data (or minimize loss of science data), then response effectiveness is measured by how much science data is acquired, or conversely, how much is lost.

Calculating response effectiveness is generally grounded in physics-based or logic-based analysis. Time plays a major role in response effectiveness. This is because the success of the mitigation is largely based on the result of a race condition of the failure effects versus the failure response. That is, failure effects take time to propagate from cause to the compromise of the relevant goal. The failure response is generally effective if it performs its function before the failure effect compromises the goal, and conversely is ineffective if it does not. For the FMCL of which the response is a part, the time to be assessed includes the entire set of detections, diagnostics, decision, and response, not just of the response by itself. The end result is an estimate of the probability that the response completes before the failure effect compromises the goal.

I. FMCL Effectiveness

Ultimately, what is needed is the ultimate metric that the FMCL as a whole successfully mitigates the failure in the given failure scenario, and then summing across all failure scenarios, provide a measure of the effectiveness of all FMCLs across all probability-weighted failure scenarios.

For any given failure scenario, the response effectiveness probability by itself is not the final result, because failure of the detections, diagnostics, or decisions also lead to the failure of the FMCL. In addition, the False Positives of the state estimates, or incorrect response decisions also create additional costs beyond simply failing to mitigate the relevant failure. All of these must be factored into the FMCL effectiveness estimate. In essence, all FN state estimates lead to the failure effect not being mitigated. All FP state estimates result in additional costs, because it activates a response action when one was not needed and hence generate new failure scenarios that did not exist before the implementation of the FMCL. Incorrect response decisions, like FP state estimates, create added cost, but they also mean that the ongoing failure is not mitigated. All of these factor into the final estimate of the effectiveness of the FMCL. The benefit of the FMCL in preserving the goal is when all FM functions work properly.

An effective FMCL design has a significant probabilistic benefit. Compared to its benefits, failures of the should be FMCL much less likely, and for FNs and incorrect response decisions, only occur if a failure occurs to begin with. FP detections, isolations, and identifications yield entirely new failures whose costs must be subtracted from the FMCL benefits. Incorrect response decisions create new costs and also yield no failure mitigation benefits.

A final issue to be considered is that for each scenario, more than one FMCL may be invoked. On SLS for aborts, it only matters which FMCL is activated first, because the response is always the same: abort. Once an abort is activated, it does not matter if another attempt to activate occurs, because the first one will execute and only one abort response can ever happen. For other systems, there can and will be other kinds of responses that can occur. If more than one FMCL is activated in a scenario, then the metrics captured must address this.

As an example, assume a spacecraft application, in which an actuator fails to function. The failure could be in the actuator, or it could be in the control electronics transmitting data between the flight computers and the actuators, or it could be the computer itself. Assume all are redundant, and all are cross strapped. One way in which FM could operate would be to detect the failure and then swap to redundant actuators, and then wait to see if actuation occurs. Assume this fails. The FM next switches to the redundant control electronics and then waits to see if actuation ensues. This fails, so the FM now switches flight computers, and finally this works. In this design, is the FM effective because ultimately it fixed the problem? If so, then the FM metrics must account for the serial nature of the responses, which operate in this case much like repair and maintenance procedures. We have described this situation before when discussing the computation of diagnostics. In essence we have on-board algorithms doing diagnostic repair procedures much as their human counterparts might do for another system. Presumably the calculation would need to be performed in a similar way, assuming that all three actions must be accounted for, and if any of them succeeds that the FM has worked correctly. Failure of the FMCL in this case would only occur if *none* of these actions succeed prior to system failure. That is, actuation is needed to achieve a goal, and if the goal is not achieved, the failure ensues.

V. Example Calculations

In this section, we will provide an example of FM metrics calculation. For the purposes of this paper, we do not intend to address all of the complexities involved with full probabilistic representations, distributions, and methods. That is not because this is unimportant, but rather that our purpose here is to show the basics of the individual metric

calculations, and how they are combined to provide estimates of the effectiveness of FMCLs, and for all FMCLs in a system as an estimate of the effectiveness of FM as a whole for a system. Attempting to describe the full statistical detail would tend to obscure what we are trying to show. However, given that all of these methods are ultimately related to the statistical benefits and risks of FM for a system to protect system goals, probabilistic methods are inherent and required. The level of sophistication needed will depend on the level and cost of risks associated with a given system, and the resources available to assess such risks.

For our example calculation, we shall assume a planetary probe application in which providing science data back to Earth is the system goal. As FM exists to protect this goal, that means that once in space, the benefit of FM is measured against potential losses of science data. For other systems of course this can vary from protection of human life, protection of profits, or other goals as appropriate.

Our example space probe goes to a planet farther away from the Sun than Earth, such as Mars or Jupiter, and will use solar panels and batteries for power generation. It has typical subsystems that provide normal spacecraft support functions: the electrical power system provides and distributes electrical power, the thermal protection system keeps temperature ranges within appropriate bounds, the science subsystems gather science data, the command and data handling system manages the mission and stores the data, the attitude control system points the spacecraft and science instruments, and the propulsion system provides thrust needed to keep attitude control and put the spacecraft on the proper trajectory. We will assume two science instruments that require precision pointing (such as cameras), and one science instrument that does not require pointing control (such as magnetometers). The Fault Management exists to protect all of these system functions, insofar as they are related to the ultimate system goal of gathering and transmitting science data to Earth.

As will be discussed below, it will be necessary to define some initial candidate FM capabilities that will be designed into the system, as the metrics will evaluate the effectiveness of those proposed designs. For our planetary probe, we will assume a dual-string design for most of the propulsion capabilities and other mechanical and moving parts, and also a centralized duplex computer system operating on a hot-warm basis (one being prime, another backup) with bus cross-strapping between the triplex and duplex components. If it is desired to evaluate different design options such as a triplex or quad computer system architecture or non-cross-strapped operations, or different variations of these, then the calculations described below can be performed for each alternative design that is being assessed.

A. Failure Scenarios and FMCLs to be Assessed

The first task is to define failure scenarios for the system. It is desirable to define all “credible” scenarios for the system, for which credibility is a probabilistic risk threshold value against which all possible scenarios are judged. Barring this, then a single failure tolerance criteria or some other relevant criteria can be used to define scenarios. A major factor in defining failure scenarios is variation in failure effects based on changes in system configuration or environment. Another key factor is to define scenarios to the level of granularity required to assess the performance of the potential FMCLs that could be designed into the system. For example, if there are three different detections and 2 different responses, which can execute in any combination, then at least 6 scenarios at the relevant level are likely needed to differentiate the performance of each of these potential combinations. It is thus necessary to have some initial idea of the FM designs that are to be assessed.

At a high level, failure scenarios relate to losses of function that can occur at different phases of the mission, and the differing failure effects and timing that are relevant in these phases. For example during science data gathering phases, loss of ability to point the science instruments is critical as it causes loss of science data immediately. However, a loss of attitude control during the cruise phase on the way to the planet causes little or no loss of science data of the target planet, as long as it is fixed prior to orbit insertion or flyby, or any other science data gathering. In general, electrical failures cause system failure within milliseconds, whereas thermal protection failures could take hours, days, or weeks. The TTCs associated with these scenarios vary accordingly. We will look at a few cases of specific FMCLs and how their effectiveness is estimated using the metrics described above. We will not attempt to define a global set of scenarios for our example space probe, but will pick out a few relevant cases to illustrate the calculation methods.

The example scenarios that will be used in our example calculations will be as follows:

- 1) Global loss of computing function during cruise phase (GLC-C)
- 2) Global loss of computing function during orbit insertion (GLC-O)
- 3) Global loss of computing function during science-gathering phase (GLC-S)
- 4) Loss of thermal control during cruise phase (LTC-C)
- 5) Loss of thermal control during science-gathering phase (LTC-S)
- 6) Loss of attitude control during cruise phase (LAC-C)

- 7) Loss of attitude control during orbit insertion (LAC-O)
- 8) Loss of attitude control during science-gathering phase (LAC-S)
- 9) Propellant leaks during cruise phase (PRP-C)
- 10) Propellant leaks during science-gathering phase (PRP-S)

There will be several or many causes that can create the conditions described in this list. For the loss of computing function, assume that a watchdog timer or some other non-computer mechanism exists to reboot the computers. Failures could be of hardware or software. Note that two scenarios, thruster leaks and loss of thermal control during orbit insertion (THR-O and LTC-O) have been removed from consideration. In our example, we have made preliminary determination that these scenarios are not credible during those particular phases, meaning that they are either improbable, or do not create mission-threatening effects during these time periods.

The probability of these scenarios are assumed estimated using standard reliability techniques from historical data and reliability handbook data, and are represented in convenient table form. Note that there are uncertainties and distributions for these values, but for simplicity of explanation of the methods, we will ignore them.

Table 1. Example Failure Scenarios and Estimated Risk Values

	Cruise (C)	Orbit Insertion (O)	Science Gathering (S)
Loss Computing (GLC)	1E-4	1E-5	3E-4
Loss of Thermal (LTC)	1E-5	Not Credible	2E-5
Loss of Att Ctl (LAC)	2E-4	3E-4	5E-4
Propellant Leak (PRP)	3E-4	Not Credible	4E-4

Total Risk in these scenarios: 2.14E-3.

The FMCLs to be assessed against these scenarios are as follows:

- 1) Watchdog Timeout; Computer Reboot and Safing (WT-Rbt&Sf)
- 2) Mission Operations Thermal Trending; Attitude Profile Redesign and Power On Extra Components (Trnd-Att&PWR)
- 3) Attitude Control Failure Detection; Sequenced Switch to Redundant Strings of Sensors, Data Buses, and Computers, and Shut Thruster Valves and Switch to Redundant Thrusters (not necessarily in that order) (ACFD-ACRED)
- 4) Excessive Imbalanced Thruster Commanding; Shut Thruster Valves and Switch to Redundant Thrusters, Sequenced Switch to Redundant Strings of Sensors, Data Buses, and Computers (not necessarily in that order, though thruster valve closures and switches would likely be the first action) (TCMD-THREDSW)

The FMCL definitions define the proposed detection, and after the semicolon the proposed response. In our FMCL design, we do not allow responses to execute in parallel, so as to avoid potential response conflicts and overly complex detection-response interactions. This rule is violated during orbit insertion, when any possible responses must be executed immediately and in parallel as the time constraints are critical. This means greater risks of the responses interacting in dangerous ways, but given that the maneuver must succeed and all possible actions to ensure success must be taken, this risk is acceptable. Also, note that failure detections, prognostics, and responses can be on board and/or on the ground. Thus the safing response after computer reboot will also include ground intervention to get the mission going again after the spacecraft goes into safing mode and attempts to contact Earth.

B. Detection Coverage

The first major metric required is failure detection is coverage. A detection mechanism can be assessed for effectiveness only against those failures, degradations, or anomalies that it theoretically capable of detecting. If it cannot detect these behaviors, even in theory, then the detection has no “coverage” of those behaviors. Conversely, if behaviors are potentially detectable by the detection mechanism, then the mechanism “covers” them. For purposes of our calculations, prognostic predictions act as failure detection mechanisms and are treated as such.

Note that not all failures are detectable even “in theory”. For loss of computer function, it could be that a total loss of power is the cause, in which case the watchdog timer will be irrelevant. One can imagine propellant leaks in the near the tanks that produce little or no rotational thrust and hence not detectable through attitude control or excessive thruster command imbalance detections. Presumably these could be detected through another mechanism such as loss of pressure in the propellant tanks, but for we are not assuming the existence of that kind of detection. Thus some small fraction of propellant leak failures will lead directly to loss of mission. During cruise this would mean no science is gathered. In the science-gathering phase, some science might have been gathered prior to occurrence of the failure.

In general, loss of science numbers during science gathering are accumulated at 50%, under the assumption that science gathering is a linear function of time, and the probability of failure is constant throughout the science gathering phase. Conversely, non-detected failures during cruise or orbit insertion would result in 0% science gathered, or 100% of the non-detected risk accruing to a loss of mission (LOM) risk. The “50%” accumulation during science phases is accounted for here under detection coverage in that an additional 50% of coverage beyond the initial coverage estimate is provided. Thus the original estimate of 99% coverage for the Watchdog Timer is improved to 99.5% for the science-gathering phase, and the original 70% estimate for the thruster command imbalance detection is increased to 75%, under the assumption that all non-covered risk (10% of the total scenario risk) is accrued to this detection instead of attitude control. This will have the effect of decreasing the loss of science in the science-gathering phase by 50% in these failure scenarios.

In our scenarios, the thermal trending prognostics and watchdog timers exclusively detect their respective scenarios, but for attitude control and leak cases, the two detection mechanisms of commanded thruster imbalance and attitude control loss can potentially detect loss of attitude control cases and thruster leak cases. After assessment of the failure effect propagations and the detection and prognostic mechanisms, we get the following coverage estimates and resulting “covered risk” and “non-covered risk” that leads to Loss of Mission because no detection can pick up the failure.

Table 2. Detection Coverage of Failure Scenarios

	WT-RbtSf	Trnd-Att&Pwr	ACFD-ACRed	TCMD-THRed
GLC-C	99%			
GLC-O	99%			
GLC-S	99.5%			
LTC-C		100%		
LTC-S		100%		
LAC-C			90%	10%
LAC-O			100%	
LAC-S			90%	10%
PRP-C			20%	70%
PRP-S			20%	75%

The next table takes the original scenario risk value, multiplies by the coverage fractions, and generates the amount of risk picked up by each detection. Any residual non-detected risk is assigned a loss of science (that is, loss of mission-LOM) number, depending on when it occurred.

Table 3: Risk Coverage and non-Covered Risk Absolute Values

	Scenario Risk	WT-RbtSf	Trnd-Att&Pwr	ACFD-ACRed	TCMD-THRed	Non-Covered LOM Risk
GLC-C	1E-4	9.9E-5				1E-6
GLC-O	1E-5	9.9E-6				1E-7
GLC-S	3E-4	2.985E-4				1.5E-6
LTC-C	1E-5		1E-5			0
LTC-S	2E-5		2E-5			0
LAC-C	2E-4			1.8E-4	2E-5	0
LAC-O	3E-4			3E-4		0
LAC-S	5E-4			4.5E-4	5E-5	0
PRP-C	3E-4			6E-5	2.1E-4	3E-5
PRP-S	4E-4			8E-5	3E-4	2E-5

C. Detection Effectiveness

In theory, even if a detection mechanism should be able to detect the relevant failure behaviors, imperfections in the detection mechanism and inherent tradeoffs in the ability to detect failure behavior from nominal behavior means that their effectiveness is always less than 100%. For detections that detect binary signals, the issues of setting thresholds to separate nominal from failed behaviors do not exist, so for these, there is one less source of errors than for detections that require the use of thresholds.

For our example detections, not all of them are effective equally at all times during the mission. The watchdog timer is not as effective during orbit insertion, because it is based on two means to reset the timer. The first is an occasional reset from the ground, and the other is an occasional reset by the flight software. The ground resets are not relevant during the short period of orbit insertions, so failures of the computing system during that time period that cause loss of most computer functions but nonetheless reset the watchdog timer will probably lead to loss of mission. Propellant leak detections are very effective in cruise, but not as effective during science-gathering, because the spacecraft maneuvers a lot during science gathering, which means that the threshold of thrust command imbalance cannot be set as tightly. More leaks will go undetected for too long. Mission operations trend analysis is assumed to be 100% effective, as there is generally plenty of time to detect a slow-acting thermal issue.

As described in the previous section regarding coverage, the effect of detection failure accrues to loss of science at only a 50% rate compared to the same failure occurring during orbit insertion or cruise phases. During these phases if the failures occur and the system is lost, no science will be gathered. The reduced effect of failures during the science-gathering phase for these cases are accounted for by increases in detection effectiveness for cruise phases for attitude control and propellant leak scenarios.

Table 4. Detection Effectiveness per Failure Scenario

	WT-RbtSf	Trnd-Att&Pwr	ACFD-ACRed	TCMD-THRed
GLC-C	99.9%			
GLC-O	95%			
GLC-S	99.9%			
LTC-C		100%		
LTC-S		100%		
LAC-C			99%	99%
LAC-O			98%	
LAC-S			99.5%	97.5%
PRP-C			99%	99%
PRP-S			99.5%	97.5%

Applying these values to the residual risk numbers from the coverage estimates, we get the following absolute residuals of detected and non-detected risk values. Note that non-detected cumulative risk adds new detection effectiveness risks to prior coverage risks. Also, for scenarios in which more than one detection picks up risk in the scenario, non-detected risks accumulate from both detections.

Table 5. Detection Effectiveness Cumulative Risks per Failure Scenario

	WT-RbtSf	Trnd-Att&Pwr	ACFD-ACRed	TCMD-THRed	Cumulative LOM Risk
GLC-C	9.89E-5				1.099E-6
GLC-O	9.405 E-6				5.95E-7
GLC-S	2.984E-4				1.649E-6
LTC-C		1E-5			0
LTC-S		2E-5			0
LAC-C			1.782E-4	1.98E-5	2E-6
LAC-O			2.94E-4		6E-6
LAC-S			4.478E-4	4.875E-5	3.5E-6
PRP-C			5.94E-5	2.079E-4	3.27E-5
PRP-S			7.96E-5	2.925E-4	2.79E-5

D. Fault Isolation Effectiveness

Once detected, the location of the cause of the failure effects must be determined, so that the proper response can be determined and executed. For slow-acting thermal problems, there is plenty of time to determine the likely location and the areas affected by it. For loss of computing, there is no other location that can cause the inability for the watchdog timer to be reset, except for the computer or the timer itself for which the risk is negligible. Problems with attitude control must be associated with the components in the control loop, with the possible exception of a debris strike. As long as the responses executed switch all possible components in the loop, these are addressed, unless it is a common-cause software problem, in which case a computer reboot may be needed. As long as this is in

the response chain, then all possibilities have been addressed. A propellant leak is the most difficult failure to isolate properly.

In our example calculation, the only cases for which isolation effectiveness will be less than 100% are for loss of control and propellant leak cases during science mission phases. For these cases, we will assume that these would lead to minor losses of science of about 1% during the mission due to this cause, or 99% effective. Thus all of the data remains the same from Table 5 above, except for the LAC-S and PRP-S scenarios. The 50% reduction factor does not apply in this case for the science-gathering phase, as the issue is not loss of the entire mission after the failure, but rather losses of science during the fault management actions, followed by the mission being re-started.

Table 6. Isolation Effectiveness Cumulative Risks per Failure Scenario

	WT-RbtSf	Trnd-Att&Pwr	ACFD-ACRed	TCMD-THRed	Non-Detected Cumulative Risk
GLC-C	9.89E-5				1.099E-6
GLC-O	9.405E-6				5.95E-7
GLC-S	2.984E-4				1.649E-6
LTC-C		1E-5			0
LTC-S		2E-5			0
LAC-C			1.782E-4	1.98E-5	2E-6
LAC-O			2.94E-4		6E-6
LAC-S			4.433E-4	4.826E-5	8.465E-6
PRP-C			5.94E-5	2.079E-4	3.27E-5
PRP-S			7.88E-5	2.89E-4	3.162E-5

E. Failure Response Decision Effectiveness

Deciding what response should be taken is generally a pre-determined activity for on-board responses, but not always for ground-based decisions. As described above, the on-board responses are generally pre-planned sequences of actions that account for all possible responses that could lead to the failure effects being observed. As long as they are all used as planned, the response decision effectiveness should be 100%. It would be less than 100% if for some reason the wrong responses were executed. This would most likely be a software or parameter load problem with a very small likelihood. Ground-based decisions are not always pre-planned and pre-analyzed in this manner. For example, slow-acting thermal problems would not be expected to occur, but there are a number of real-life examples in which this has occurred. Deciding the best response to take is not always simple, and could in fact be found inadequate when executed.

For our purposes here, the lack of a proper on-board response will be addressed under response decision effectiveness, using the reasoning that sometimes an effective a response could in theory have been put in place. It will apply to the ground-based thermal failure responses during the science mission phase. Here we will postulate that a response that changes the mission profile and configuration to yield better thermal control do not work as well as expected, and another response is selected instead, but only after some science data is compromised.

We will assume a 1% loss of science data could occur for thermal problems due to improper responses taken by mission operations during science-gathering phase only. This yields the following residual risks.

Table 7. Response Decision Effectiveness Cumulative Risks per Failure Scenario

	WT-RbtSf	Trnd-Att&Pwr	ACFD-ACRed	TCMD-THRed	Cumulative LOM Risk
GLC-C	9.89E-5				1.099E-6
GLC-O	9.405E-6				5.95E-7
GLC-S	2.984E-4				1.649E-6
LTC-C		1E-5			0
LTC-S		1.98E-5			2E-7
LAC-C			1.782E-4	1.98E-5	2E-6
LAC-O			2.94E-4		6E-6
LAC-S			4.433E-4	4.826E-5	8.465E-6
PRP-C			5.94E-5	2.079E-4	3.27E-5
PRP-S			7.88E-5	2.896E-4	3.162E-5

F. Failure Response Effectiveness

The last metric for the FMCLs is failure response effectiveness. Like the prior metrics, there is a response effectiveness fraction for each FMCL in each scenario, which is multiplied by the cumulative reduced risk fraction from prior steps, in this case the cumulative risk from the failure response decision function. Any portion of that risk for which the response is not effective is added to the cumulative LOM risk. Thus if a response is 90% effective, 90% of the risk resulting from prior steps leads to a successful response, and hence to a successful outcome for the FMCL as a whole. The 10% that is not effective accumulates to the cumulative LOM risk for that scenario.

Failure response effectiveness is the portion of the calculation in which the overall FMCL latency is addressed. That is, the time it takes to execute the FMCL is compared to the Time to Criticality associated with the failure effects in the scenario. If the TTC time expires before the FMCL completes, then the response is not effective. Otherwise, the response generally is effective. As the TTC and the FMCL latency are both probabilistic predictions of their respective times, this leads to a probabilistic effectiveness criteria. Other factors can also play a part in the estimate of response effectiveness, such as variations in the criticality of the failure effects, and not just in the timing of those failure effects. We will not attempt to address all of these possibilities here in this paper, but suffice it to say that like all other calculations described, physical and probabilistic factors must be addressed to perform these estimates, which vary based on the application.

For our example case, prognostics and ground response to thermal failure is generally performed well in advance of the TTC, but the nature of the response could be such that one inherently compromises some of the mission data gathering. This can occur if one must keep the spacecraft at specific attitudes to keep certain parts in view of the Sun's rays, and keeping other parts shaded. This could take away from times in which it would otherwise be desirable to point the science instruments. This contributes to a prediction of less than 100% effectiveness of responses during the science gathering phase of the mission.

Responses to computer problems result in computer reboot and safing will cause loss of science data during science-gathering. A loss of computer function during orbit insertion causes loss of mission every time as a reboot response will cause loss of control due to the length of time required for a reboot compared to the timing of a loss of control during this time period, leading to the orbit insertion not occurring. Loss of attitude control can cause loss of science data during science gathering as the instruments would not be pointed at their targets for some period of time, though most likely this would be for smaller time periods than if a computer reboot and safing were required. Propellant loss cases would be less likely to cause loss of science, except insofar as loss of propellant shortens the mission. The following table shows the estimates of response effectiveness for our example case.

Table 8. Response Effectiveness per FMCL per Failure Scenario

	WT-RbtSf	Trnd-Att&Pwr	ACFD-ACRed	TCMD-THRed
GLC-C	95%			
GLC-O	0%			
GLC-S	92%			
LTC-C		96%		
LTC-S		98%		
LAC-C			100%	98%
LAC-O			90%	
LAC-S			99%	95%
PRP-C			98%	98%
PRP-S			98%	95%

As with previous FMCL functions, the response effectiveness numbers are multiplied by the risk residuals from the prior table for response decision effectiveness to yield the total FMCL effectiveness and total cumulative LOM risk per scenario in the table below.

Table 9. FMCL Cumulative Risks per Failure Scenario

	WT-RbtSf	Trnd-Att&Pwr	ACFD-ACRed	TCMD-THRed	Cumulative LOM Risk
GLC-C	9.396E-5				6.044E-6
GLC-O	0				1E-5

GLC-S	2.745E-4				2.522E-5
LTC-C		9.6E-6			4E-7
LTC-S		1.94E-5			5.96E-7
LAC-C			1.782E-4	1.94E-5	2.396E-6
LAC-O			2.646E-4		3.54E-5
LAC-S			4.388E-4	4.585E-5	1.531E-5
PRP-C			5.821E-5	2.037E-4	3.805E-5
PRP-S			7.723E-5	2.751E-4	4.768E-5

G. Value of FMCLs

The value of the FMCLs per scenario can now be gathered up to estimate the value of the FMCLs across the scenarios, which is to say the value that each FMCL provides for the system in terms of protecting the mission, which in this case means protection of science return. One simply adds up the risk that is mitigated by each FMCL from all scenarios in which the FMCL applies. After removing the amount of risk in each scenario that is not covered, that is to say not detectable even in theory, one can divide the total risk mitigated by the FMCL by the amount of covered risk in the scenario by the FMCL to get the effectiveness fraction or percentage of the FMCL. The non-covered fraction should be reported separately as “non-covered risk” that creates loss of mission / loss of science. If it is desired to protect against that non-covered risk, then some other mechanism must be provided, either FM to predict or detect and respond, or to prevent the failure from occurring by adding design margin (which is the other major aspect of SHM), which would directly lower the probability of occurrence.⁷

This yields the following value estimates of the FMCLs in terms of loss of science / loss of mission. The positive way of stating the value is “Loss of Science Benefit”, or “Loss of Mission Benefit”.

Table 10. FMCL Value Estimates

FMCL Name	Original Risk	Non-Covered Risk	Covered Risk	LOS/LOM Benefit	Effectiveness Fraction
WT-RbtSf	4.10E-4	2.60E-6	4.07E-4	3.68E-4	90.4%
Trnd-Att&Pwr	3.00E-5	0	3.00E-5	2.90E-5	96.7%
ACFD-ACRed	1.07E-3	0	1.07E-3	1.02E-3	95.1%
TCMD-THRed	6.30E-4	5.00E-5	5.80E-4	5.44E-4	93.8%

For our example case, we see that the most effective FMCL is Thermal Trending with Attitude-Power response. However, it is the lowest value (lowest LOS/LOM Benefit) FMCL because the probability of thermal problems is low. However, because this is done on the ground, it is likely low cost compared to on-board options and thus probably worthwhile. Others are less effective against the failures they address, but are nonetheless high overall value because the probability of failure scenarios that these address is high. FM designers can use metrics like this to determine whether the value of the FMCLs are worth their cost in terms of hardware redundancy, algorithms, testing, and so on. For example, going from a relatively ineffective dual computer architecture to a much more effective triplex would definitely decrease system risks, but would require additional mass, power, and cost for the third computer.

FM metrics like this, and the design considerations addressed in calculating them and using the results to aid in design decisions, have been used and demonstrated successfully on the NASA Space Launch System program. On SLS, the application was the selection of algorithms to detect abort conditions, which are crew-threatening failures and from which the crew needs to escape and return to Earth using launch and on-orbit abort systems.⁸

VI. Conclusion

Fault Management theory is based on the idea that Fault Management operates as a suite of meta-control loops (FMCLs) that predict, detect and respond to existing or prospective failures, and keep or return the system to a controllable state. The purpose of these control loops is to provide insurance and assurance that the system will achieve its goals. Whether these FMCLs are worth the cost and complexity of adding them into the design is a design trade like any other. As with most other engineering trades, quantification is a great help in determining if the design is technically effective, and cost effective.

Control theory suggests the use of metrics based on state estimation and state control. State estimation refers to the knowledge of system health, and state control refers to the control of the system to maintain health or mitigate

failure effects to achieve as many of the system's goals as possible. Control of system health will not be effective unless the system's health is properly estimated. State estimation metrics are based on True Positive / True Negative / False Positive / False Negative measures. State control metrics are based on decision and control effectiveness, which for the latter are largely dependent on the timing race condition between failure effect propagation and the latency of FMCL execution.

This paper describes the theoretical issues that must be considered, and provides an example calculation to demonstrate how to use these metrics to determine the value of Fault Management Control Loops. This in turn can be used to assess trade studies on different FM designs, and in the latter part of the design and build process to help verify and validate that design to determine if system dependability goals are being met. That this can be done is demonstrated by its successful use in NASA's human-rated Space Launch System for assessment of abort detection and response.

Acknowledgments

Some of the work that led to this paper was performed under NASA Small Business Innovative Research contracts S5 05 NNX15CM31P, and NNX16CM10C to Qualtech Systems, Inc.

Stephen B. Johnson thanks John C. Day for his extremely valuable collaboration on the development of the theory of System Health Management and Fault Management, as noted in the series of Johnson-Day papers on the subject from 2010-2014 as noted in the reference list below. Many of the details of the theory were worked out in practice on the NASA Space Launch System Program, with Yunnhon Lo in particular being critical to developing the specific calculation methods and relationships to probabilistic risk assessment methods. Jonathan Breckenridge was crucial in the implementation of these techniques. Both were co-authors in the 2014 Lo, Johnson, and Breckenridge paper that provided an overview of these calculations, listed in the references.

References

- ¹⁶Johnson, S. B., "The Theory of System Health Management," Chapter 1 in *System Health Management: with Aerospace Applications*, edited by Johnson, S. B., Gormley, T. J., Kessler, S. S., Mott, C., Patterson-Hine, A., Reichard, K. M., and Scandura, Jr., P. A., John Wiley United Kingdom, Chichester, United Kingdom, 2011, pp. 3-27.
- ²Johnson, S. B., "Conceptual Framework for a Fault Management Design Methodology," *AIAA Infotech@Aerospace Conference 2011*, AIAA-2010-227006, Atlanta, Georgia, April 2010.
- ³Johnson, S. B., and Day, J. C., "System Health Management Theory and Design Strategies," *AIAA Infotech@Aerospace Conference 2011*, AIAA-2011-977233, St. Louis, Missouri, March 29-31, 2011.
- ⁴Johnson, S. B., "Goal-Function Tree Modeling for Systems Engineering and Fault Management," *AIAA Infotech@Aerospace Conference 2013*, AIAA-2013-4576, Boston, Massachusetts, August 19-22, 2013.
- ⁵Melcher, K. J., Cruz, J. A., Johnson, S. B., and Lo, Y., "Abort Trigger False Positive and False Negative Analysis Methodology for Threshold-Based Abort Detection," *Prognostics and Health Management Society 2015 Annual Conference*, Coronado, California, October 18-24, 2015.
- ⁶Roemer, M. J., Byington, C. S., Kacprzynski, G. J., Vachtsevanos, G., and Goebel, K., "Prognostics," Chapter 17 in *System Health Management: with Aerospace Applications*, edited by Johnson, S. B., Gormley, T. J., Kessler, S. S., Mott, C., Patterson-Hine, A., Reichard, K. M., and Scandura, Jr., P. A., John Wiley United Kingdom, Chichester, United Kingdom, 2011, pp. 281-295.
- ⁷Day, J. C., and Johnson, S. B., "Fault Management Design Strategies," *Space Operations 2014 Conference*, Pasadena, California, May 2014.
- ⁸Lo, Y., Johnson, S. B., and Breckenridge, J., "Application of Fault Management Theory to the Quantitative Selection of a Launch Vehicle Abort Trigger Suite," *IEEE Prognostics and Health Management Conference*, Spokane, Washington, June 2014.